arXiv:1207.3202v1 [cs.CG] 13 Jul 2012

Embedding the dual complex of hyper-rectangular partitions

Michael Kerber*

Abstract

A rectangular partition is the partition of an (axis-aligned) rectangle into interior-disjoint rectangles. We ask whether a rectangular partition permits a "nice" drawing of its dual, that is, a straight-line embedding of it such that each dual vertex is placed into the rectangle that it represents. We show that deciding whether such a drawing exists is NP-complete. Moreover, we consider the drawing where a vertex is placed in the center of the represented rectangle and consider sufficient conditions for this drawing to be nice. This question is studied both in the plane and for the higher-dimensional generalization of rectangular partitions.

1 Introduction

Problem statement and results We consider the partition of the d-dimensional cube into axis-parallel d-dimensional hyper-rectangles (or boxes) with disjoint interiors which we call *hyper-rectangular partitions*. The dual complex of such a partition is an abstract simplicial complex of dimension d and represents the connectivity of the boxes in the partition, possibly after a slight distortion if more than $d+1$ boxes meet in a common point. Each vertex of the dual represents a box of the partition.

We pose the question of whether a partition permits a "nice" drawing of its dual in $\mathbb{R}^d$. We impose three conditions for niceness: first of all, the drawing must be straight, that is, a face must be drawn as the convex hull of the vertices that are on its boundary. That means that the drawing is completely determined by the positions of the vertices. Second, we disallow improper intersections of faces; for instance, edges are not allowed to cross except at their endpoints. The first two conditions can be summarized to that the drawing must yield a simplicial complex in $\mathbb{R}^d$. The final condition is that each vertex must be placed in the box that it represents; that guarantees some minimal amount of geometric closeness between a primary cell of the partition and its dual counterpart. If a drawing with these three conditions is possible, we call it an *embedding*. See Figure 1 for an illustration of a rectangular partition in $\mathbb{R}^2$ and its dual.

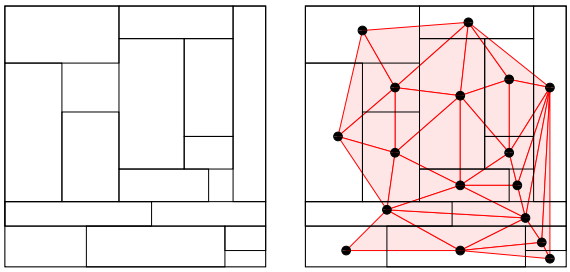


Figure 1: Left: A (generic) rectangular partition. Right: An embedding of the dual.

We will see that not every hyper-rectangular partition permits an embedding of the dual with the properties from above, so a natural question to ask is how fast can the existence of an embedding be decided for a given partition. Our first main contribution is to show that this decision problem is NP-hard already in the planar case. More precisely, we restrict to the case that the boxes of a partition and the vertices of the dual are snapped to a uniform grid, and we prove that the problem is NP-complete under this constraint; see Section 3 for the precise statement.

The most natural choice for placing a vertex of the dual (for arbitrary dimension d) is the center of the corresponding box. We call a partition *center-embeddable* if this vertex placement extends to an embedding of the dual. We ask for simple sufficient conditions on a partition to be center-embeddable besides the trivial condition that all boxes are d-dimensional cubes of the same size. We call a partition β-balanced (with $\beta \geq 1$) if for any set of k boxes intersecting in a common point, the ratio of the longest side among all k boxes divided by the shortest side of all k boxes is bounded

*IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria.

by β. As our second main contribution, we investigate the relationships between β-balanced and center-embeddable partitions: In the planar case, we show that center embeddability is guaranteed if the partition is $(3-\epsilon)$-balanced for any $\epsilon > 0$. However, in $\mathbb{R}^3$ (and higher dimensions) we can construct a $(1+\epsilon)$-balanced partition that is not center-embeddable. The situation changes if we restrict to *cubical partitions*, consisting only of d-dimensional cubes: in this case, $(3-\epsilon)$-balanced partitions are center-embeddable in $\mathbb{R}^3$, and we can construct a cubical β-balanced partition in $\mathbb{R}^d$ that is not center-embeddable for any $\beta > d/(d-2)$. It remains open whether cubical β-balanced partitions with $\beta < d/(d-2)$ are center-embeddable in dimensions larger than 3.

Related work Rectangular partitions in $\mathbb{R}^2$ have a lot of applications, for instance in VLSI design [11], cartography [13], and database-related application [12]; we refer to the survey by Eppstein [7] for more examples. In this context, rectangular partitions sometimes appear as the *rectangular dual* of a triangulated planar graph. Linear-time algorithms have been presented for the computation of such a rectangular partition dual [10]. Our situation, however, is different as the rectangular partition is the input object and we ask for a straight-line embedding of its dual where the vertices are constrained to lie inside rectangles. This can be seen as an instance of planar graph embedding with constraints; NP-hardness has been shown for other constraints, such as when fixing the length of each edge [4], and, more related to our approach, when restricting the placement of each vertex to a disc [9]. Furthermore, the problem of simultaneously embedding a planar graph and its dual on the integer grid received some interest, e.g. [8]. Most related to our approach is the variant where the embedding of the primal graph is fixed and an embedding of the dual is seeked for such that and primal-dual pairs of edges intersect. Bern and Gilbert [2] show that the problem is linear-time solvable if all faces are convex and four-sided, and becomes NP-hard for convex five-sided faces. The latter is proven with a reduction from planar 3SAT and is similar to the proof presented in this work. However, although rectangular partitions can be seen as planar graphs with convex faces, there is no direct reduction from our problem because we allow the dual graph to intersect the partition arbitrarily, not just at primal-dual pairs.

The higher dimensional equivalent of rectangular partitions (and their dual complexes) apparently have not been investigated from a theoretical point of view. Our motivation for this topic is originated in the approximation and simplification of d-dimensional image data [1]. The idea is to identify rectangular regions in which the image looks "similar". In the simplest case, similarity means that the image values inside a region are similar, but different measures can be defined. The regions define a hyper-rectangular partition which can consist of substantially less elements then the original n^d image cells. On the other hand, turning a (hyper-)rectangular structure into a simplicial structure is a common operation in geometry processing; see [14] for a discussion with references. The standard approach is to triangulate the rectangular regions separately in a conformal way; this however, runs into difficulties in our setup with different face lengths and intersection patterns and generally leads to a large number of simplices. The dual complex construction avoids these drawbacks, to the price that it does not necessarily embed in the ambient space nicely.

In [1], the special case of images in $\mathbb{R}^3$ induced by an oct-tree data-structure has been investigated. In that work, the dual complex is used to approximate the persistent homology [6] of the image. It was also shown that the dual complex is center-embeddable, provided that the underlying oct-tree is balanced. In [5], that result is generalized to the higher-dimensional analog of oct-trees, called *hierarchical cubical subdivisions*; Balanced subdivisions are special cases of 2-balanced partitions in our notation. Our results show that the hierarchical structure is crucial for obtaining the general result of [5], as for all $d \geq 5$, we can construct a non-hierarchical 2-balanced partition that is not center-embeddable.

Outline We introduce hyper-rectangular partitions and their dual complexes formally in Section 2. Section 3 is devoted to the NP-completeness proof of finding an embedding in the planar case. We study center-embeddability in Section 4. Section 5 concludes the paper.

2 Hyper-Rectangular Partitions and Dual Complexes

This section introduces the most important concepts needed for the results of this work.

We call a point set of the form $[a_1, b_1] \times \ldots \times [a_d, b_d]$ with $a_i < b_i$ and $a_1, \ldots, a_d, b_1, \ldots, b_d \in \mathbb{Z}$ an *integral hyper-rectangle*, or just a *box*, with *lengths* $b_1 - a_1, \ldots, b_d - a_d$. An *integral hyper-square*, or *square box* is a box where all lengths are equal. We can think of a box to be composed out of *hyper-pixels*, or *unit boxes* which are integer translates of $[0,1]^d$. Let $B = [0,n]^d$ be a square box with an arbitrary $n > 0$. A *hyper-rectangular partition* $C = (R_1, \ldots, R_m)$ (with $m \leq n^d$) of B is a collection of boxes R_i such that their union equals B and their interiors are disjoint. If the union of the R_i is only a subset of B, we call the collection a *partial hyper-rectangular partition*. We usually omit the term "hyper-rectangular" and just talk about a *partition*. A partition is called *generic* if not more than $d+1$ boxes intersect in a common point.

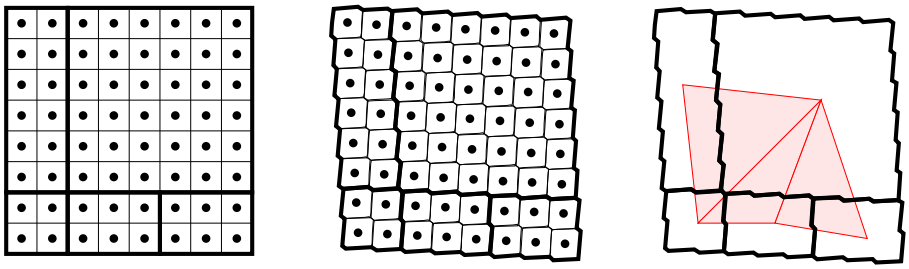


Figure 2: Left: A non-generic rectangular partition with 5 rectangles. Middle: The distorted pixels and rectangles. Right: The dual complex of the partition.

We associate an abstract d-dimensional *dual complex* $D(C)$ to a partition C. For a generic partition, $D(C)$ is just the nerve of C, that is, each box R_i is represented by a vertex v_i, and the simplices of $D(C)$ correspond to subsets of $\{R_1, \dots, R_m\}$ with non-empty intersection. This construction fails for non-generic partitions because intersections of $d+2$ rectangles give rise to a $(d+1)$-simplex in the nerve. To circumvent this problem, we slightly perturb the boxes to obtain a generic situation, following the construction from [5]: The union of unit boxes can be seen as the Voronoi diagram of the pixel centers $P := \{(x_1 + 1/2, \dots, x_d + 1/2) \mid x_1, \dots, x_d \in \mathbb{Z}\}$. For an arbitrary $\epsilon \in (0, 1)$, we define the *distortion* of a point $p = (x_1, \dots, x_d) \in P$ as

$$T_\epsilon p := (x_1 - \epsilon\frac{\Sigma x_i}{2}, \dots, x_d - \epsilon\frac{\Sigma x_i}{2}).$$

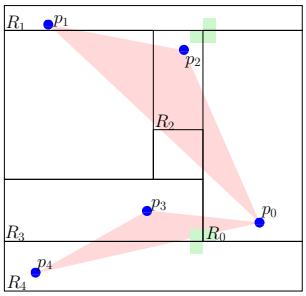
We consider the Voronoi diagram of the distorted pixel centers. For a pixel U, we define its distortion $\tilde{U}$ as the Voronoi cell of its distorted pixel center. For a box consisting of pixels $U_1, \dots, U_k$, we define its distortion as the union of the distorted pixels $\tilde{U_1}, \dots, \tilde{U_k}$. Finally, we define the dual complex $D(C)$ of a partition C to be the nerve of the distorted boxes in the partition. See Figure 2 for an example in the plane. We remark that the dual complex does not depend on the choice of ϵ. Informally, the distortion is a way to remove high-dimensional simplices from the nerve in non-generic situations, thereby "preferring" connections between hyper-rectangle in the diagonal direction $(1, \dots, 1)$; we refer to [5] for more details.

For a hyper-rectangular partition $C = (R_1, \dots, R_k)$, let v_i denote the vertex of the dual complex $D(C)$ that represents R_i. We call a mapping from $\{v_1, \dots, v_k\}$ to $\mathbb{R}^d$ a *projection*. A projection is *integral* if each v_i is mapped to a point $(a_1/2, \dots, a_d/2)$ with $a_1, \dots, a_d \in \mathbb{Z}$; in other words, v_i is mapped to the first barycentric refinement of the integer grid. A projection is *faithful* if each v_i is mapped in the interior of R_i. We assume in this work that projections are integral and faithful, unless otherwise stated. A projection π extends to the whole dual complex $D(C)$ by mapping a higher-dimensional simplex to the interior of the convex hull of the projected boundary vertices. Abusing notation, we let π also denote the extended mapping. We call a projection an *embedding*, if this mapping is injective, or in other words, if the image is a simplicial complex in $\mathbb{R}^d$.

Seed configurations Consider a d-simplex σ of $D(C)$, dual to the intersection point w of the boxes $R_0, \dots, R_d$. By definition of $D(C)$, there is a *unique* collection of unit boxes $U_0, \dots, U_d$ with $U_i \subseteq R_i$ such that the distorted unit boxes $\tilde{U_0}, \dots, \tilde{U_d}$ intersect as well. By definition, the U_i intersect in w. Let u_i be the center of U_i. We call $(u_0, \dots, u_d)$ the *seed configuration* of σ.

A sequence of points $p_0 = (x_{0,1}, \dots, x_{0,d}), \dots, p_d = (x_{d,1}, \dots, x_{d,d})$ in $\mathbb{R}^d$ has the *orientation*

$$O(p_0, \dots, p_d) = \operatorname{sign} \det \begin{pmatrix} 1 & x_{0,1} & \dots & x_{0,d} \\ \vdots & & \ddots & \\ 1 & x_{d,1} & \dots & x_{d,d} \end{pmatrix}.$$



We say that a projection π of $D(C)$ to $\mathbb{R}^d$ *preserves the orientation of the d-simplex σ*, if $O(\pi(v_0), \dots, \pi(v_d)) = O(u_0, \dots, u_d)$, where $u_0, \dots, u_d$ is the seed configuration of σ. For an equivalent formulation in the plane, assume that R_0, R_1, R_2 are arranged counterclockwisely around w. Then π preserves the orientation of σ, if the circular path $\pi(v_0), \pi(v_1), \pi(v_2)$ is counterclockwisely arranged, too. In the picture on the right, the seed configuration of the dual 2-simplices $\{R_0, R_1, R_2\}$ and $\{R_0, R_3, R_4\}$ is illustrated; note that the given projection does not preserve the orientation of the former 2-simplex, but does preserve the orientation of the latter.

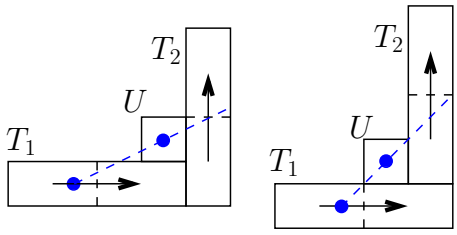


Figure 3: The two possible configurations of L-joints (up to rotations and reflections).

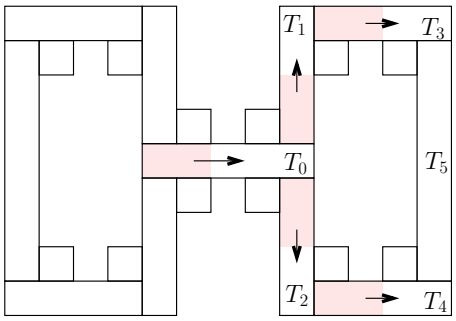


Figure 4: A (partial) partition that does not permit an embedding

The following result is a generalization of the "Geometric Realization Theorem" from [5]. In there, it was proven that for balanced hierarchical cubical subdivisions, the projection that maps a vertex to the center of the cube is always an embedding. The main property exploited in the proof is the "Orientation Lemma", stating that the orientation of each d-simplex is preserved for such subdivisions. This is no longer true for arbitrary partitions, but the same proof can be used to show (Appendix B gives more details):

EMBEDDING THEOREM. Let C be a partition. A projection π of $D(C)$ is an embedding if and only if it preserves the orientation of each d-simplex.

The Embedding Theorem gives a *local* criterion for the existence of an embedding. All our results are eventually reduced to investigate the orientation of projected d-simplices. Because of the importance of the theorem in our work, we included Appendix B which repeats the proof argument from [5] for the convenience of the reader.

3 NP-completeness

In this section, we concentrate on the planar case $d = 2$. We adapt our notation and leave out the prefix "hyper-" on all defined terms.

Thin rectangles and L-joints We call an integral rectangle *thin* if one of its side has length 1, and the other side has length at least 4. The *length* ℓ of a thin rectangle is the length of its longer side. Clearly, any projection has to map the vertex of a thin rectangle to one of the ℓ pixel centers. Three intersecting rectangles are called an *L-joint* if two of the rectangles are thin and their union is L-shaped (up to rotation and reflection), and the third rectangle is the pixel in the bulge of the L-shape; see Figure 3. If the L-joint belongs to a partition, it induces a 2-simplex in $D(C)$. We call a concatenation of L-joints *L-path*. Such an L-path can be cyclic, in this case we call it *L-cycle*. It is convenient to think of the thin rectangles in an L-joint being directed consistently towards one end of the L-shape. If the direction is fixed, it makes sense to talk about the *front rectangle* and *back rectangle* of the L-joint, about the *front half* and *back half* of each thin rectangle, and about the *front pixel* and the *back pixel* which are the extremal pixels in the a thin rectangle.

L-JOINT LEMMA. Let (T_1, T_2, U) be a directed L-joint, with T_1 being the back rectangle, T_2 being the front rectangle, and U being the bulge pixel. Let π be a projection that preserves the orientation of the 2-simplex induced by the L-joint. Then, if π maps the vertex of T_1 to to its back half, it also maps the vertex of T_2 to its back half.

PROOF. Figure 3 displays the two possible connections of T_1 and T_2, if both are of length 4. Recall that π is assumed to be integral and, therefore, has only 4 choices to place the vertex for both T_1 and T_2. We can see that, if π preserves the orientation and maps the vertex of T_1 into its back half it must map the vertex of T_2 below the dashed line, so this vertex must go into the back half in both cases. The situation does not change if we extend T_1 to the left or T_2 to the top. □

We can use L-joints to show that dual complexes cannot always be embedded: Consider the (partial) rectangular partition in Figure 4. Let π be any projection of $D(C)$. Assume for a contradiction that π preserves the orientation of all 2-simplices induced by the partial partition. Assume w.l.o.g. that the vertex of T_0 is projected to the left half of T_0 (otherwise, the symmetric argument applies). Direct the thin rectangles according to the arrows in the figure. By the L-joint lemma, the vertices of T_1 and T_2 must be placed in their back halves, and by repeating the argument, the same holds for the rectangles T_3 and T_4. Using the L-joint Lemma again, the vertex of T_5 must be placed in the upper half of T_5 (caused by the L-joint with T_3) and also in the lower half of T_5 (caused by the L-joint with T_4), a contradiction. Filling out the partial partition with pixels, we obtain a full partition C such that no projection preserve the orientation of all

triangles. The Embedding Theorem asserts that there is no embedding.

The reduction We define the decision problem `Faithful_Embed` as follows: for a partition $C = (R_0, \ldots, R_m)$ of $[0, n] \times [0, n]$, is there an embedding of $D(C)$? Our goal is to show NP-completeness of this problem.

First of all, the problem is clearly in NP: given an integral partition C, we can compute $D(C)$ and check whether a specific projection causes an orientation switch for any triangle in polynomial time.

For the reduction, we define the `grid3sat` problem in the same way as in [9]:

> `grid3sat`: Given some $N \times N$ grid (with N linear in n) in which some grid points are labelled as clauses and some as variables. Variables are connected with clauses by vertex-disjoint paths on the grid. A sign is associated to every such path indicating whether the corresponding variable is negated in the clause or not. Every clause is incident to exactly three paths. Is the formula described in this way satisfiable or not? See Figure 5 (left) for an example.

As stated in [9], `grid3sat` is NP-complete. Consider an instance I of `grid3sat` and let F be encoded formula (in 3-CNF form). Based on the grid graph, we will construct a rectangular partition C such that the existence of a satisfying assignment for F is equivalent to the existence of an embedding of C.

We barycentrically subdivide the grid given by I a constant number of times (in fact, 5 barycentric subdivisions suffice) and construct our partition inside this base grid. For that, we first replace vertices representing variables or clauses by boxes of sufficient size (which we fill later), and we replace paths connecting variables and clauses by disjoint L-paths. Each path inherits the sign of the corresponding path in I. We think of L-paths being directed from variables to clauses. See Figure 5 (right).

Inside boxes representing a variable X_i, we place a **variable gadget** (Figure 6 (left)). In its center, it has a L-cycle with 4 thin rectangles of length 8, which we direct in clockwise direction. We call this cycle the *variable cycle for* X_i. Each L-path to a clause box that starts at a boundary side of the gadget is extended into the interior of the box and forms a L-joint with the corresponding side of the variable cycle. If the path is associated with a plus sign (that is, the variable appears in non-negated form in the corresponding clause), the path is connected to the front pixel. If it is associated with a minus sign, it is connected to the back pixel.

Inside boxes representing a clause, we place a **clause gadget** (Figure 6 (right)). The three L-paths arriving at the boundary of the gadget are extended to connect to the *clause rectangle*, with the last rectangle being exactly 4 base squares long. If the L-paths arrive at different sides than displayed in Figure 6, we can just rotate the gadget by a multiple of $90°$. The gadget is designed to satisfy the following two properties:

CLAUSE LEMMA. Let v_0, v_1, v_2, v_3 be the vertices of R_0, R_1, R_2, R_3 as in Figure 6 (right), and let π be a partial projection that maps v_i to p_i for $1 \leq i \leq 3$.

1. If p_1, p_2, and p_3 are in the back half of R_1, R_2, and R_3, respectively, then π cannot be completed to an embedding.
2. If at least one of the p_i is in the front pixel of R_i, then there exists an extension of π to v_0 such that π preserves the orientation of all triangles incident to v_0.

PROOF. For the first part, let p_0 be the projection of v_0 in a completion of π. If π is an embedding, p_0 is left of ℓ_1, right of ℓ_2 and below ℓ_3, which is impossible. For the second part, if p_1 or p_2 are in the front pixel, we can set p_0 to be the center of the rightmost or leftmost dashed base square in Figure 6 (right), respectively. If p_3 is in the front pixel, we can choose the center of the topmost blue base square for p_0. From the two red lines in Figure 6 (right), it can be seen that the triangles indeed have the correct orientation. □

After all, we have constructed a partial integral rectangular partition, consisting of variable cycles, clause rectangles, and L-paths connecting them. We finally complete the partition by filling the empty spots with pixels. Let C denote this partition. We can construct C in polynomial time in n. The following lemma is sufficient to prove NP-hardness of `Faithful_Embed`.

LEMMA. There exists an embedding of $D(C)$ if and only if F has a satisfying assignment.

PROOF. "$\Rightarrow$": Let π be an embedding. We define an assignment as follows: for every variable X_i, consider its variable cycle (which is clockwisely arranged). Note that the L-joint Lemma implies that π either projects all vertices of the variable cycle to the front half, or it projects all vertices to the back half. We set X_i to 1 if π projects to the front half and to 0 otherwise. We show that this assignment satisfies F:

Let $C = L_i \vee L_j \vee L_k$ be a clause of F where $L_i = X_i$ or $L_i = \bar{X}_i$, and consider the corresponding clause gadget. Let R_i, R_j, R_k the last elements of the L-paths connecting the clause to the variable gadgets of X_i, X_j, and X_k, respectively. Since π is an embedding, at least one of the vertices

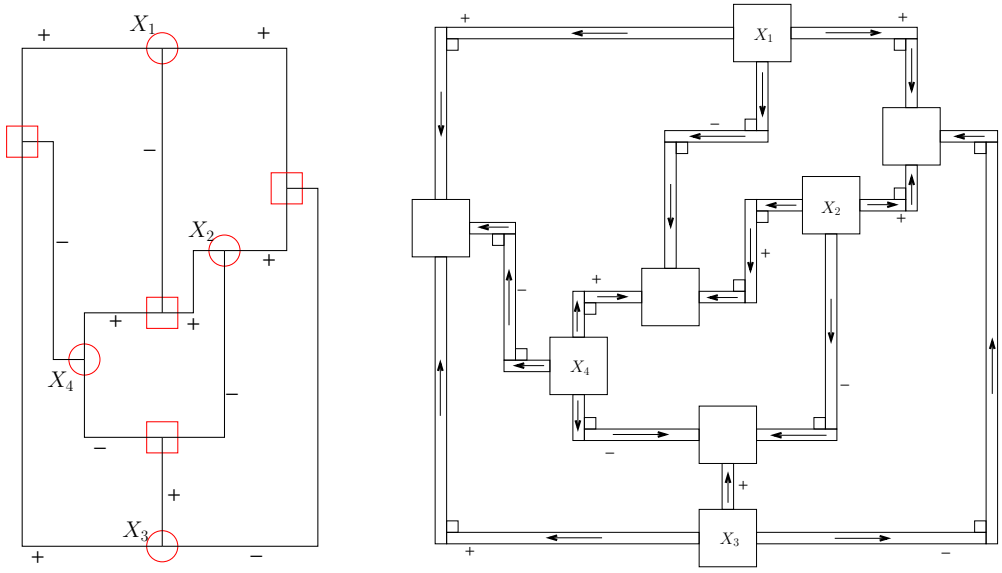


Figure 5: Left: A graph presenting the formula $(X_1 \vee X_2 \vee \bar{X}_3) \wedge (\bar{X}_1 \vee X_2 \vee X_4) \wedge (X_1 \vee \bar{X}_3 \vee \bar{X}_4) \wedge (\bar{X}_2 \vee X_3 \vee \bar{X}_4)$. Right: A corresponding representation as a partial integral rectangular partition, with L-path connections.

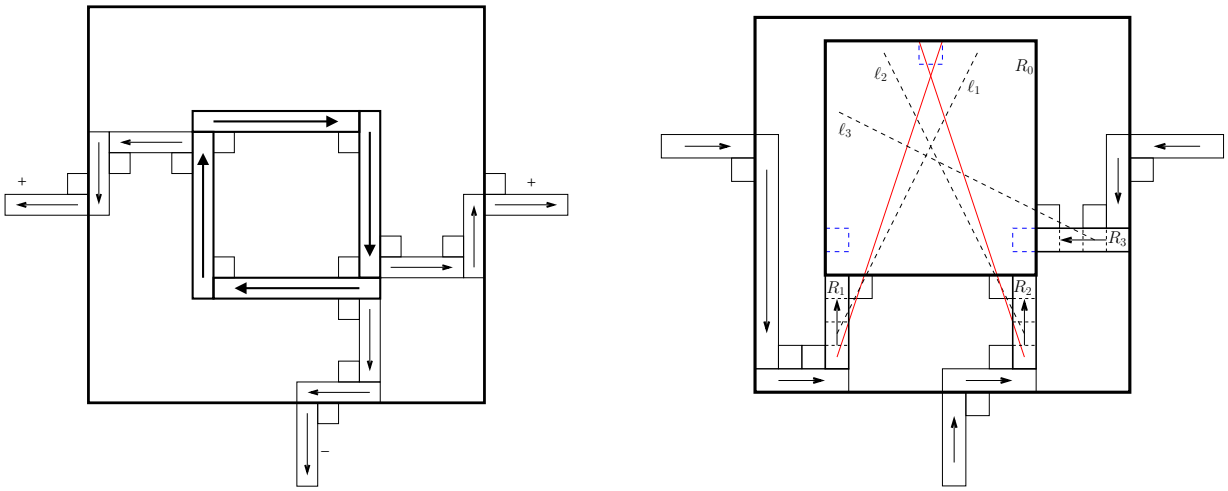


Figure 6: Left: The variable gadget for X_1 from Figure 5. Right: The clause gadget for $\bar{X}_2 \vee X_3 \vee \bar{X}_4$ from Figure 5.

of R_i, R_j, or R_k is placed in the corresponding front half by the Clause Lemma. Assume w.l.o.g. that the vertex of R_i is placed in the front half. By iteratively applying the L-joint Lemma, all vertices in the L-path from X_i to C must therefore be placed in the front half. If $L_i = X_i$, the path was associated with a plus sign, and therefore, the L-path is connected with the variable cycle of X_i at a front pixel (as for instance the left and right paths originating from the variable cycle in Figure 6 (left)). Since π preserves orientations, it follows that the vertex of the variable cycle must also have been placed in the front half. Hence, our assignment sets X_i to 1 and therefore, the clause is satisfied. If $L_i = \bar{X}_i$, the argument is analogous.

"$\Leftarrow$": We construct a projection based on a satisfying assignment. We start with the variable gadgets: If $X_i = 1$, we place the vertices in its variable cycle at the front pixel, otherwise at the back pixel. For every L-path that leaves a variable cycle, we distinguish two cases: if the path has a plus sign and $X_i = 1$, or if the path has a minus sign and $X_i = 0$, we place all vertices in the path in the front pixel; otherwise, we place all vertices in the back square. It is straight-forward to verify that this projection preserves the orientation for each 2-simplex caused by L-joints. Finally, let $C = L_i \vee L_j \vee L_k$ be a clause as above. Consider the corresponding clause gadget and let R_i, R_j, R_k denote the connecting (thin) L-path rectangles. By our construction, the vertex of R_i is in the front pixel if and only if $L_i = 1$, same for j and K. Since we have a satisfying assignment, at least one literal is satisfied, so one of the vertices is in the front pixel. By the Clause Lemma, we can therefore place the vertex of the clause rectangle such that all triangles preserve their orientation.

After all, our projection preserves all orientations of the constructed partial partition. It is not difficult to see that the additional 2-simplices caused by filling the empty spots with base squares preserve their orientation as well (in fact, any projection preserves their orientation). Therefore, the constructed projection is an embedding by the Embedding Theorem. ◻

After all, we can summarize

THEOREM. `Faithful_Embed` is NP-complete.

4 Center projections and β-balancing

We turn back to arbitrary dimensions. We define the *center* of a box $[a_1, b_1] \times \ldots \times [a_d, b_d]$ to be the point $((a_1 + b_1)/2, \ldots, (a_d + b_d)/2)$. The *center projection* is the projection of $D(C)$ that maps each vertex to the center of the corresponding box. Clearly, this projection is integral and faithful. We call C *center-embeddable* if the center projection is an embedding.

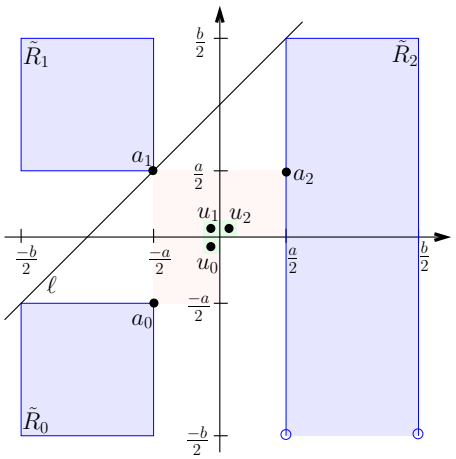


Figure 7: Illustrations of the proof of Theorem 6 for $\beta = 3$

We define the *balance* of a set of boxes $\{R_0, \ldots, R_k\}$ with non-empty intersection to be the length of the longest side among all boxes $R_0, \ldots, R_k$ divided by the length of the shortest side among all boxes $R_0, \ldots, R_k$. We define the balance of a k-simplex σ of $D(C)$ to be the balance of the set of dual boxes. The *aspect ratio* of a box R is the balance of $\{R\}$. Obviously R has aspect ratio 1 if and only if R is a square box. We call $D(C)$ β-*balanced* for some $\beta \geq 1$, if each simplex of $D(C)$ has a balance of at most β. Informally, β-balanced partitions have the property that boxes are not too skinny, but also neighboring boxes do not differ too much in side lengths.

The simplest sufficient condition for a partition to be center-embeddable is that all boxes are square boxes of same size. We investigate several generalizations of this trivial criterion.

Planar results It is not hard to prove that a partition in $\mathbb{R}^2$ consisting only of squares is always center-embeddable, because the edge connecting the centers of two squares does not leave the union of the squares. However, if we bound the aspect ratio of any rectangle by $1 + \epsilon$ (with $\epsilon > 0$), this property does not hold in general, and we can construct a counterexample for center-embeddability. For β-balanced partitions, we prove a tight bound for center-embeddability.

PLANAR BALANCING THEOREM. A β-balanced partition of $\mathbb{R}^2$ is center-embeddable for $\beta < 3$, and there exists a 3-balanced partition that is not center-embeddable.

PROOF. Consider a β-balanced partition. For the first part, it is enough to prove that the center projection preserves the orientation of each 2-simplex for $\beta < 3$ according to the Embedding Theorem. So, let σ be a 2-simplex, caused by the intersection of three rectangles R_0, R_1, R_2. W.l.o.g., we

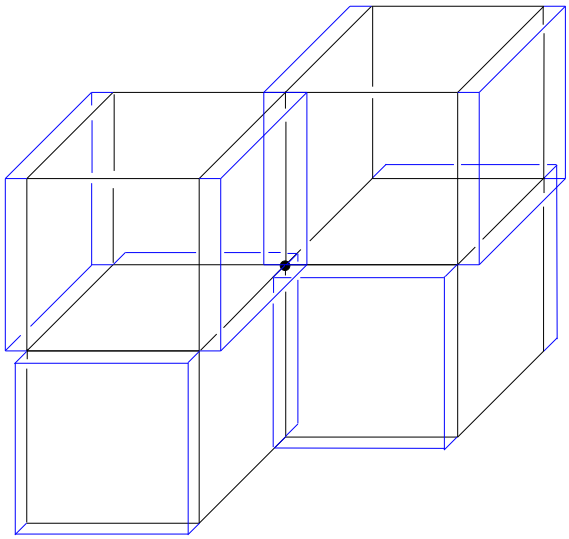


Figure 8: Illustrations of the proof of Theorem 7 (right).

assume that the intersection point is the origin, that R_0 is completely contained in the lower-left quadrant, R_1 is contained in the upper-left quadrant, and R_2 is contained in the halfplane $x \geq 0$; all other cases are obtained by suitable rotations and reflections. Let a denote the length of the shortest side among the R_i, and let b denote the length of the longest side; $\frac{b}{a} \leq \beta$. We let $\tilde{R}_i$ denote the set of possible positions of the center of R_i. By the imposed constraints, it is straight-forward to verify that $\tilde{R}_i$ are rectangles as illustrated in Figure 7 (left). Note that the lower side of $\tilde{R}_2$ does not belong to the set of possible centers.

Recall that the orientation of σ is preserved if $O(p_0, p_1, p_2) = O(u_0, u_1, u_2)$, where the u_i define the seed configuration of σ; in our case, $u_0 = (-1/2, -1/2)$, $u_1 = (-1/2, 1/2)$, and $u_2 = (1/2, 1/2)$. We scale each u_i by a, obtaining the points a_0, a_1, a_2; obviously, $O(u_0, u_1, u_2) = O(a_0, a_1, a_2)$. Assume that the orientation of σ is not preserved. Then, if we continuously move the points a_i to the points p_i along the connecting line segment (which lies in $\tilde{R}_i$, there must be a time where the three points become colinear. It follows that there must be a line in $\mathbb{R}^2$ which intersects each $\tilde{R}_i$. Such a line, however, exists if and only if $b \geq 3a$, as can be seen by elementary geometric arguments. It follows that the orientation of σ is preserved if $\beta < 3$.

The construction also yields the counterexample for $\beta = 3$: Consider the rectangles intersecting in the origin with centers $p_0 = (-3/2, -1/2)$, $p_1 = (-1/2, 1/2)$, $p_2 = (1/2, 3/2)$. We can fill the bounding box by pixels to obtain a 3-balanced partition. The 2-simplex dual to the origin is projected to a line segment because p_0, p_1, p_2 are colinear. Therefore, the partition is not center-embeddable. □

Higher dimensions Restricting to square boxes (that is, aspect ratio 1) does not guarantee center-embeddability for any $d > 2$: In $\mathbb{R}^3$, the triangle spanned by the centers of three intersecting cubes might leave the union of the three cubes, and a counterexample can be constructed; see [1] and [5]. We show that also restricting to β-balanced partitions does not guarantee embeddability, unless $\beta = 1$:

THEOREM. For any $\beta > 1$, there exists a β-balanced partition in $\mathbb{R}^d$ that is not center-embeddable.

PROOF. We restrict to $d = 3$ for the simplicity of presentation although our construction generalizes to any dimension: Choose $b \in \mathbb{Z}$ such that $\beta \geq \frac{b+2}{b} > 1$. Consider four cubes of length b having the origin as a corner, such that the centers lie in a common plane (Figure 8 (right)). Then, extend the cubes by a layer of voxels at two opposite faces, such that all voxels around the origin are covered by the four boxes. This ensures that the 3-simplex induced by the four boxes belongs to the dual. The extension does not change the centers, consequently, the center projection cannot preserve the orientation of that simplex. Moreover, all boxes have side length b and $b + 2$. It remains to show that the arrangement of the four boxes can be completed to a partition without destroying β-balancing. We refer to Appendix A for details of this technical step. □

The cubical case We finally consider β-balanced partitions in $\mathbb{R}^d$ where all boxes are square boxes. We call such partitions *cubical*. We first construct counterexamples in arbitrary dimension to show:

THEOREM. For $d \geq 3$ and $\beta > \frac{d}{d-2}$, there exists a β-balanced cubical partition in $\mathbb{R}^d$ that is not center-embeddable.

PROOF. For $d \geq 3$, $a, b \in \mathbb{Z}$ positive, and $\delta \in \{0, 1\}$ consider the $(d+1) \times (d+1)$-matrix

$$M_d^{(\delta)} := \begin{pmatrix} 1 & -\frac{a}{2} & \dots & & -\frac{a}{2} \\ 1 & \frac{b}{2} & -\frac{b}{2} & \dots & -\frac{b}{2} \\ 1 & -\frac{b}{2}+\delta & \frac{b}{2} & -\frac{b}{2} & \dots & -\frac{b}{2} \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 1 & -\frac{b}{2}+\delta & \dots & -\frac{b}{2}+\delta & \frac{b}{2} & -\frac{b}{2} \\ 1 & -\frac{b}{2}+\delta & \dots & \dots & -\frac{b}{2}+\delta & \frac{b}{2} \end{pmatrix}$$

It can be shown by elementary row operations (Appendix C) that

$$\det M_d^{(0)} = \frac{1}{2} b^{d-1}(da - (d-2)b).$$

In particular, note that the determinant is positive if $\frac{b}{a} < \frac{d}{d-2}$, vanishes if $\frac{b}{a} = \frac{d}{d-2}$, and negative if $\frac{b}{a} > \frac{d}{d-2}$.

Let a and b be such that $\beta > \frac{b}{a} > \frac{d}{d-2}$, so that $\det M_d^{(0)} < 0$. If we choose a and b large enough, it follows that $\det M_d^{(1)} < 0$ as well, just because $M_d^{(1)}$ constitutes a small perturbation of $M_d^{(0)}$ for large a and b. The rows of $M_d^{(1)}$ define $d+1$ points $p_0, \ldots, p_d$ in $\mathbb{R}^d$ (ignoring the first column). We define the square box R_i with center p_i and side length a for R_0, and length b for $R_1, \ldots, R_d$. By this choice, it can be seen easily that all 2^d unit boxes around the origin are covered by the R_i, so $R_0, \ldots, R_d$ form a d-simplex σ in the dual complex. Moreover, $u_0 := (-1/2, \ldots, -1/2) \in R_0$, $u_1 := (1/2, -1/2, \ldots, -1/2) \in R_1$, $u_2 := (1/2, 1/2, -1/2, \ldots, -1/2) \in R_2$, and so on, thus $u_0, \ldots, u_d$ form a seed configuration and it is easily seen to that $O(u_0, \ldots, u_d) > 0$. Because $O(p_0, \ldots, p_d) < 0$ by construction, the center projection does not preserve the orientation of σ.

The last step of the proof is to show that we can complete the initial configuration $R_0, \ldots, R_d$ to a complete β-balanced cubical partition in $\mathbb{R}^d$. See Appendix A for details. □

We are able to show that the constructed counterexample is the worst case in $\mathbb{R}^3$. Precisely, we state:

THEOREM. Cubical β-balanced partitions in $\mathbb{R}^3$ are center-embeddable for $\beta < 3$.

PROOF. For simplicity, we restrict to generic partitions in the proof; the general case works with the same methodology, but requires closer investigation of the distortion defined for the dual complex construction. As in Theorem 7, we use the terms cube and voxel instead of square box and hyper-pixel.

Consider $d+1$ cubes $R_0, \ldots, R_d$ intersecting in a point p. Let a be the length of the shortest and b be the length of the longest cube among the $R_0, \ldots, R_3$. Assume w.l.o.g. that $p = 0$. By genericity, the union of the R_i covers all voxels adjacent to p, and each cube covers at least one voxel. Represent these voxels by the 3-bit numbers $0, \ldots, 7$ where the ℓ-th bit of voxel j is set to 1 if and only if the x_ℓ-coordinate of that voxel is positive. Assign to each R_i a subset of $\{0, \ldots, 7\}$, denoting the voxels that it occupies. Of course, only three cases are possible, namely that R_i covers a single voxel, that R_i covers two voxels which are face adjacent, or four voxels which lie in a common halfplane. We consider C_i, the set of possible centers for R_i. If R_i covers only one voxel adjacent to the origin, e.g., voxel 0, then C_i is the line segment connecting $(-a/2, -a/2, -a/2)$ and $(-b/2, -b/2, -b/2)$. If R_i covers two voxels which are face adjacent, say voxels 0 and 1, C_i is the trapezoid spanned by the four points $(\pm a/2, -a/2, -a/2)$, $(\pm b/2, -b/2, -b/2)$; more precisely, the center can not lie on either of the two non-parallel lines because this would prevent the cube to span over both voxels, so C_i is the trapezoid with those two sides excluded. Finally, if R_i covers four voxels in a common hyperplane, say voxels $4, \ldots, 7$, C_i is the polytope spanned by the 8 points $(\pm a/2, \pm a/2, a/2)$, $(\pm b/2, \pm b/2, b/2)$; the polytope has 6 faces, two of them being parallel. More precisely, the center can not lie on a boundary face except for the two parallel sides, so C_i is that polytope with the other 4 faces excluded.

Assume for a contradiction that the center projection does not preserve the orientation of the simplex spanned by $R_0, \ldots, R_3$. Scaling the seed configuration of the simplex by a factor of a, we obtain a simplex spanned by $q_0, \ldots, q_3$ with q_i in the closure of C_i with the same orientation as the seed configuration. Because this orientation is not preserved, there are points $p_0, \ldots, p_3$ with $p_i \in C_i$ such that the orientation is different. It follows that there are points $p'_0, \ldots, p'_3$ with $p'_i \in C_i$ whose orientation is zero. In other words, there exists a plane which intersects $C_0, \ldots, C_3$. However, $C_0, \ldots, C_3$ are completely determined by the way of how the boxes R_i decompose the voxels $\{0, \ldots, 7\}$ adjacent to the origin, and there are only two configurations possible, up to rotations and reflections: the *regular configuration*

$$\{0\}, \{1\}, \{2, 3\}, \{4, 5, 6, 7\},$$

and the *singular configuration*

$$\{0, 1\}, \{2, 3\}, \{4, 6\}, \{5, 7\}.$$

Consequently, there must be a plane intersecting the two line segments, the trapezoid and the polytope defined by the regular configuration, or a plane intersecting the four trapezoids defined by the singular configuration. It can be verified, however, that such a plane does not exist by formulating the statement in terms of a quantified system of inequalities and using a quantifier elimination program [3][1] (see Appendix D). Hence, the center-projection preserves the orientation of every simplex and is an embedding. □

The proof idea can in principle be extended to higher dimensions. However, there are two problems: first, the number of configurations to check increases; for instance, there are already 3 configurations to check in $\mathbb{R}^4$. Second, and more seriously, the complexity of the quantifier elimination increases dramatically for higher dimensions; in fact, we were not able to verify that β-balanced cubical partitions in $\mathbb{R}^4$ are center-embeddable for $\beta < \frac{d}{d-2} = 2$.

[1] We used qepcad: `http://www.usna.edu/cs/~qepcad/B/QEPCAD.html`

5 Conclusion

We have proven positive and negative results about partitions and their dual in this paper; it is NP-hard to decide whether an embedding exists already in the plane; if we project vertices to centers of rectangles, there are some simple balancing conditions that ensure embeddability in the plane, namely, if all rectangles are squares or if the partition is $(3-\epsilon)$-balanced with $\epsilon > 0$. In higher dimensions, however, we need to combine both properties, that means, consider cubical $(3-\epsilon)$-balanced partitions to get a guarantee that non-trivial partitions can be embedded. We are posing the question whether another simple condition could guarantee center-embeddability as well.

Quad-tree decompositions and their higher-dimensional analogues have the nice property that they can be balanced with a greedy strategy, without granulating the partition too much. It seems unclear, however, how a planar partition can be turned into a 3-balanced partition efficiently by subdividing rectangles.

This work has considered hyper-rectangular partitions in general – can better embeddability results be derived for certain restrictions, in particular for partitions that arise from a sequence of hyperplane cuts (sometimes also called guillotine constructions)? We expect the answer to be negative in the plane; however, the situation might be different in higher dimension.

On a technical side, a natural improvement would be to replace the computer-assisted proof in Theorem 9 with a geometric argument; this would hopefully extend into higher dimensions and prove our conjecture that β-balanced partitions in $\mathbb{R}^d$ are center-embeddable for $\beta < \frac{d}{d-2}$. A second question is about a variant of `Faithful_embed`: while it is clear that deciding the existence of an embedding remains hard in three dimensions, it is not clear whether the restriction to cubical partitions simplifies the problem.

We have concentrated on integral partitions and projections in this work because of the sake of simplicity, and because the situation is not too artificial, for instance in the context of image processing. The non-integral situation is not much different: by a similar, but more technical proof, it can be shown that `Faithful_embed` remains NP-hard, but it is unclear whether the problem remains in NP. Also, our results on β-balanced partitions are the same for non-integral partitions.

References

[1] P. Bendich, H. Edelsbrunner, M. Kerber: "Computing Robustness and Persistence for Images". *IEEE Transactions on Visualization and Computer Graphics* **16** (2010) 1251–1260.

[2] M. Bern, J. R. Gilbert: "Drawing the planar dual". *Inf. Process. Lett.* **43** (1992) 7–13.

[3] G. E. Collins, H. Hong: "Partial Cylindrical Algebraic Decomposition for quantifier elimination". *Journal of Symbolic Computation* **12** (1991) 299–328.

[4] P. Eades, N. Wormald: "Fixed Edge Length Graph Drawing is NP-hard". *Discrete Applied Mathematics* **28** (1990) 111–134.

[5] H. Edelsbrunner, M. Kerber: "Dual Complexes of Cubical Subdivisions in $\mathbb{R}^n$". *Discrete and Computational Geometry* **47** (2012) 393–414.

[6] H. Edelsbrunner, D. Letscher, A. Zomorodian: "Topological Persistence and Simplification". *Discrete and Computational Geometry* **28** (2002) 511–533.

[7] D. Eppstein: "Regular Labelings and Geometric Structures". *CoRR* **abs/1007.0221** (2010).

[8] C. Erten, S. G. Kobourov: "Simultaneous Embedding of a Planar Graph and Its Dual on the Grid". *Theor. Comp. Sys.* **38** (2005) 313–327.

[9] M. Godau: "On the difficulty of embedding planar graphs with inaccuracies". In: *Graph Drawing, LNCS*, vol. 894, 254–261, 1995.

[10] X. He: "On Finding the Rectangular Duals of Planar Triangular Graphs". *SIAM J. Comput.* **22** (1993) 1218–1226.

[11] C.-C. Liao, H.-I. Lu, H.-C. Yen: "Compact floorplanning via orderly spanning trees". *Journal of Algorithms* **48** (2003) 441 – 451.

[12] S. Muthukrishnan, V. Poosala, T. Suel: "On Rectangular Partitionings in Two Dimensions: Algorithms, Complexity and Applications". In: *Database Theory — ICDT'99, LNCS*, vol. 1540, 236–256, 1999.

[13] E. Reisz: "The Rectangular Statistical Cartogram". *Geographical Review* **24** (1934) 292–296.

[14] K. Weiss, L. De Floriani: "Diamond Hierarchies of Arbitrary Dimension". *Computer Graphics Forum* **28** (2009) 1289–1300.

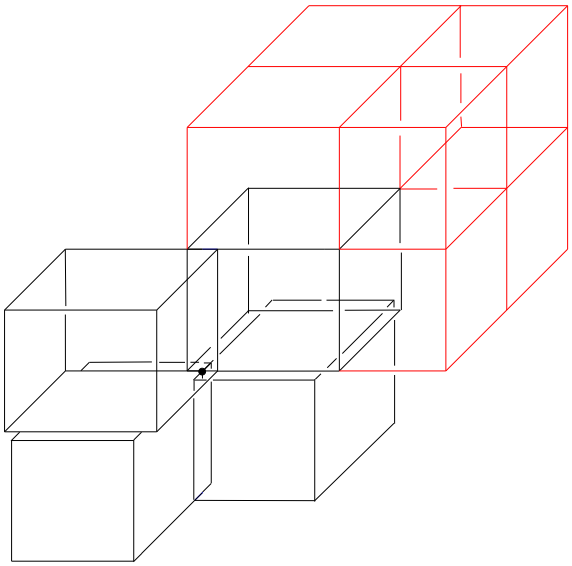


Figure 9: A part of the complete partition, after the split of one of the $\tilde{B}_i$.

A Details of the counterexample constructions

We have created configurations in Theorems 7 and 8 which lead to a simplex whose orientation is not preserved under the center projection. We have to show that those initial configurations can be completed to a β-balanced partition. For Theorem 7, this is relatively straight-forward. For the proof, we need the following definition: Let C be a partition in $\mathbb{R}^d$, B a box in the partition with corners $(q_1, \ldots, q_D)$ (with $D = 2^d$), and p a point in the interior of B. *Splitting* B at p means to replace B in C with D boxes defined by the corners p and q_i, with $1 \leq i \leq D$.

LEMMA. The partial partition $R_0, \ldots, R_3$ from Theorem 7 can be completed to a β-balanced partition in $\mathbb{R}^3$.

PROOF. Let $b \in \mathbb{Z}$ as in Theorem 7, such that $\beta > \frac{b+2}{b-2}$. We will construct a partition of $B := [-2b, 2b]^3$, consisting of exactly $64 = 8 \cdot 8$ boxes, with $R_0, \ldots, R_3$ among them. Fix a corner a_i of B (with $1 \leq i \leq 8$). It lies in an octant of $\mathbb{R}^3$ with respect to the origin, and there is exactly one box R_j (with $0 \leq j \leq 3$) that extends into this octant. We consider the unique box $\tilde{B}_i$ with (opposite) corners a_i and c_j, the previously defined corner of R_j (note that c_j does not necessarily lie in the same octant). Observe that the side lengths of $\tilde{B}_i$ are $2b-1$, $2b$, or $2b+1$. Also, $B_1, \ldots, B_8$ form a partition of B, and each R_j (with $0 \leq j \leq 3$) is fully contained in some B_i.

We split each $\tilde{B}_i$ from above into 8 boxes: if $\tilde{B}_i$ contains some R_j (with $0 \leq j \leq 3$), we let c_j^* denote the corner of R_j opposite to c_j, and split $\tilde{B}_i$ at c_j; clearly, one of the 8 created boxes is R_j. See Figure 9 for an illustration. If $\tilde{B}_i$ does not contain any R_j, we simply split at an integer point closest to its center. It is simple to verify that all boxes created by these operations have side lengths of at least $b-2$ and at most $b+2$. Since $\frac{b+2}{b-2}$ by assumption, the rectangular partition is β-balanced, but not center-embeddable, because it contains $R_0, \ldots, R_3$. □

For Theorem 6, the completion of the initial configuration is technically more challenging than in Theorem 8 because we can only use square boxes to fill the empty space. Indeed, our construction needs much more filling space to "balance out" the differences between the square boxes, and it requires square boxes of many different size, all close to b.

The construction needs two preparatory lemmas. The first is a basic fact from elementary number theory

LEMMA. Let $z_0 := 1$ and $z_1, \ldots, z_k$ be the first k prime numbers (starting with $z_1 = 2$). Let $b \in \mathbb{Z}$ such that $b = \lambda \cdot (z_1 \cdots z_k) + 1$ with some $\lambda \in \mathbb{Z}$. Then, the $(k+1)$ numbers $b + z_0 - 1, \ldots, b + z_k - 1$ are pairwise coprime.

PROOF. Assume for a contradiction that $b + z_i - 1$ and $b + z_j - 1$ have a common prime factor p. Then $p = z_\ell$ for $1 \leq \ell \leq k$, because the two numbers differ by less than z_k. Moreover, $b - 1$ is a multiple of z_ℓ. It follows that z_ℓ divides both z_i and z_j, which is a contradiction. □

The next lemma proves that we can fill arbitrary large boxes in $\mathbb{R}^d$ just using square boxes with 2^d pairwise coprime side lengths.

LEMMA. Let $D := 2^d$ and let $b_1 < \ldots < b_D$ be pairwise coprime integers. Then there exists some $L_0 \in \mathbb{Z}$ such that every box with smallest side length at least L_0 can be partitioned into square boxes of side lengths $b_1, \ldots, b_D$. Moreover, the partition can be chosen such that a box of side length b_1 is in one of its corners.

PROOF. Set $B := \{b_1, \ldots, b_D\}$. Choose L_0 such that for any integer $L \geq L_0$, and any two disjoint subsets B_1, B_2 of B, L can be represented as

$$L = \lambda_1 \prod_{b \in B_1} b + \lambda_2 \prod_{b \in B_2} b$$

with λ_1, λ_2 non-negative numbers. We prove inductively that any box with minimal side length at least L_0 can be partitioned. If $d = 1$, the statement is trivial. Let B be such a box in $\mathbb{R}^d$ with $d > 1$ and let $\tilde{B}$ denote the same box projected into $\mathbb{R}^{d-1}$. Using the induction hypothesis with the numbers $b_1, \ldots, b_{D/2}$, we can partition $\tilde{B}$ with square boxes in $\mathbb{R}^{d-1}$. If we carry this partition into $\mathbb{R}^d$ with the lower side

of every box being in the plane $x_d = 0$, this does not yield a box because the square boxes have different heights in x_d-direction; however, we can stack up boxes to match their heights. Formally, define $P_1 = b_1 \cdots b_{D/2}$ and place P_1/b_i vertical copies for a square box of length b_i. In this way, we construct a box whose projection to $\mathbb{R}^{d-1}$ is $\tilde{B}$, and whose length in x_d-direction is P_1. We can do exactly the same construction using the integers $b_{D/2+1}, \ldots, b_D$, which yields a box of length $P_2 = b_{D/2+1} \cdots b_D$ in x_d-direction. Let L_d be the length of B in x_d-direction. By assumption, we can find non-negative integers λ_1, λ_2 with $\lambda_1 P_1 + \lambda_2 P_2 = L_d$. Thus, placing λ_1 copies of the box of first type and λ_2 copies of the box of second type gives the desired box. ◻

LEMMA. The partial partition $R_0, \ldots, R_d$ from Theorem 6 can be completed to a β-balanced cubical partition in $\mathbb{R}^3$.

PROOF. The first step is to choose a and b appropriately. Note that the proof of Theorem 6 requires that $\beta > \frac{b}{a} \geq \frac{d}{d-2}$, and such that $\det M_d^{(1)}$ is negative. We impose additional conditions: Let $D := 2^d$ and $z_1, \ldots, z_{D-1}$ be the sequence of the first $D-1$ prime numbers (note that the magnitude of z_{D-1} does only depend on d) We choose a and b such that $b = \lambda z_1 \cdots z_{D-1} + 1$ for some $\lambda \in \mathbb{Z}$, such that

$$\beta > \frac{b + z_{D-1} - 1}{a} > \frac{b}{a} \geq \frac{d}{d-2},$$

and such that $\det M_d^{(1)}$ is negative. In particular, applying Lemma 11 yields D pairwise coprime integers $b := b_1 < \ldots < b_D = b + z_{D-1} - 1$.

Lemma 12 applied on the sequence $b_1, \ldots, b_D$ asserts the existence of some L_0 such that any box with side lengths at least L_0 can be filled. We choose some $L \geq L_0 + 1$ that is a multiple of a, and set $B := [-L, L]^d$. We have to show that we can construct a partition of B containing $R_0, \ldots, R_d$ with square boxes of sizes $a, b_1, b_2, \ldots, b_D$ only. By our choice of a and b, the so-obtained cubical partition is β-balanced, and the claim is proven.

The construction works similar as in Lemma 10. We first partition B into D boxes $\tilde{B}_j$ (which are not square boxes in general), each anchored at one corner of B, and with $d+1$ of them containing one of the initial square boxes R_i. Each $\tilde{B}_j$ has side lengths of at least $L - 1 \geq L_0$; therefore, it can be filled with squares boxes of lengths $b_1, \ldots, b_D$ according to Lemma 12. If $\tilde{B}_j$ contains some R_i, we can enforce that R_i is in the corner of the filling, as required. A special case is that $\tilde{B}_j$ contains R_0 which is of length a. However, by construction, this $\tilde{B}_j$ is a square box of length L, and we can simply fill it with square boxes of length a because L is a multiple of a. ◻

B The Embedding Theorem

We give some more details about the proof of the Embedding Theorem (Theorem 1) which is based on [5].

For a partition $C = (R_1, \ldots, R_m)$, let $\tilde{R}_1, \ldots, \tilde{R}_m$ denote the corresponding distorted boxes as defined in Section 2.

LEMMA. Any intersection of k distorted boxes of C is either empty or an $(n-k)$-ball. In particular, any non-empty intersection is contractible.

PROOF. The proof given in the "Fractual Distortion Lemma" in [5] extends to arbitrary partitions (it is neither exploited that the boxes are cubes, nor that they are arranged hierarchically). ◻

We now give a proof of the Embedding Theorem which works similar to the proof of the "Geometric Realization Theorem" from [5].

Let $B \subset \mathbb{R}^d$ be the box that is partitioned by C. We attach two layers of unit boxes at the boundary of B, obtaining an extended box $\bar{B}$. Note that this yields a partition $\bar{C}$, which is an extension of C. We extend the projection π of C to $\bar{\pi}$ of $\bar{C}$ by mapping vertices in the new layers to the corresponding box centers.

We compactify $\mathbb{R}^d$ to the d-dimensional sphere $\mathbb{S}^d$ by adding a vertex at infinity. Similarly, we compactify the dual complex $D(\bar{C})$ by adding a new vertex at infinity and connecting it to every simplex of the boundary of $D(\bar{C})$. Because of the preceding lemma, we can apply the Nerve Theorem to $D(\bar{C})$, which states that $D(\bar{C})$ triangulates a d-dimensional ball, and therefore, its compactification triangulates $\mathbb{S}^d$. Hence, the extended projection $\bar{\pi}$ defines a continuous mapping $g : \mathbb{S}^d \to \mathbb{S}^d$. The *degree of* g at a point x not in any $(d-1)$-simplex is the number of d-simplices that contain $g^{-1}(x)$, counting a d-simplex positive or negative depending on the orientation of its image under g. This degree is 1 in between the two outside layers, because all boxes are of unit size and therefore, the orientation of all d-simplices can be chosen positive. However, the degree of a map is a global property that does not depend on the specific location of x. Hence, it is 1 for any x.

Assume now that π preserves the orientation of all d-simplices of $D(C)$. We can assume that the orientation of all seed configurations is positive; thus the orientation of all d-simplices is positive. Because the degree of g is 1, it follows that every point is in the interior of exactly one projected d-simplex. This implies that π is an embedding, proving one direction. The other direction can be proven just by reversing the argument.

C Determinant bound

We prove in this section that for the $(d+1)\times(d+1)$ matrix

$$M_d^{(0)} := \begin{pmatrix} 1 & -\frac{a}{2} & & \dots & & -\frac{a}{2} \\ 1 & \frac{b}{2} & -\frac{b}{2} & \dots & \dots & -\frac{b}{2} \\ 1 & -\frac{b}{2} & \frac{b}{2} & -\frac{b}{2} & \dots & -\frac{b}{2} \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 1 & -\frac{b}{2} & \dots & -\frac{b}{2} & \frac{b}{2} & -\frac{b}{2} \\ 1 & -\frac{b}{2} & \dots & \dots & -\frac{b}{2} & \frac{b}{2} \end{pmatrix},$$

it holds that

$$\det M_d^{(0)} = \frac{1}{2}b^{d-1}a - \frac{1}{2}(d-2)b^d = \frac{1}{2}b^{d-1}(da-(d-2)b).$$

Indeed, we subtract the $(i-1)$-st row from the i-th row for $i = n+1,\dots,2$ to obtain

$$\begin{aligned} \det M_d^{(0)} &= \det \begin{pmatrix} 1 & -\frac{a}{2} & & \dots & & -\frac{a}{2} \\ & \frac{a+b}{2} & \frac{a-b}{2} & \dots & \dots & \frac{a-b}{2} \\ 0 & -b & b & & & \\ 0 & & -b & b & & \\ \vdots & & & \ddots & \ddots & \\ 0 & & & & -b & b \end{pmatrix} \\ &= \det \begin{pmatrix} \frac{a+b}{2} & \frac{a-b}{2} & \dots & \dots & \frac{a-b}{2} \\ -b & b & & & \\ & -b & b & & \\ & & \ddots & \ddots & \\ & & & -b & b \end{pmatrix}, \end{aligned}$$

where empty spots are zero. We simplify further by factoring out $1/2$ in the second row, factoring out b in rows $3,\dots,d+1$, and shifting the second row to the bottom, thereby changing the sign of every row that goes up. Finally, we can remove the first row and column of the matrix. With these steps, we obtain

$$\begin{aligned} \det M_d^{(0)} &= \det \begin{pmatrix} \frac{a+b}{2} & \frac{a-b}{2} & \dots & \dots & \frac{a-b}{2} \\ -b & b & & & \\ & -b & b & & \\ & & \ddots & \ddots & \\ & & & -b & b \end{pmatrix} \\ &= \frac{1}{2}b^{d-1}\det \begin{pmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \\ a+b & a-b & \dots & \dots & a-b \end{pmatrix}. \end{aligned}$$

It is straight-forward to verify that the determinant of the rightmost $(d\times d)$-matrix equals the sum of its last row, which is $da-(d-2)b$.

D Details on quantifier elimination

The proof of Theorem 9 relies on the non-existence of planes that intersect certain point sets in $\mathbb{R}^3$. We will give some details on how we prove this non-existence. Recall from the proof that we defined a and b to be the shortest and longest sides among the cubes $R_0,\dots,R_3$. By scaling, we can just assume that $a=1$. Because of the assumption, it follows that $b \le \beta < 3$.

Regular case The regular configuration

$$\{0\},\{1\},\{2,3\},\{4,5,6,7\},$$

induces the point sets

$$\begin{aligned} C_0^{(b)} &:= \mathrm{hull}\{(-1,-1,-1),(-b,-b,-b)\} \\ C_1^{(b)} &:= \mathrm{hull}\{(1,-1,-1),(b,-b,-b)\} \\ C_2^{(b)} &:= \mathrm{hull}\{(-1,1,-1),(-b,b,-b),(1,1,-1),(b,b,-b)\} \\ &\quad -\mathrm{hull}\{(-1,1,-1),(-b,b,-b)\} \\ &\quad -\mathrm{hull}\{(1,1,-1),(b,b,-b)\} \\ C_3^{(b)} &:= \mathrm{hull}\{(-1,-1,1),(-b,-b,b),(-1,1,1),(-b,b,b), \\ &\qquad (1,-1,1),(b,-b,b),(1,1,1),(b,b,b)\} \\ &\quad -\mathrm{hull}\{(-1,-1,1),(-b,-b,b),(-1,1,1),(-b,b,b)\} \\ &\quad -\mathrm{hull}\{(-1,-1,1),(-b,-b,b),(1,-1,1),(b,-b,b)\} \\ &\quad -\mathrm{hull}\{(-1,1,1),(-b,b,b),(1,1,1),(b,b,b)\} \\ &\quad -\mathrm{hull}\{(1,-1,1),(b,-b,b),(1,1,1),(b,b,b)\}, \end{aligned}$$

where hull stands for the convex hull of a point set. Assume for a contradiction that e is a plane intersecting $C_0^{(b)},\dots,C_3^{(b)}$. Since $C_i^{(b)} \subset C_i^{(3)}$, e intersects $C_0^{(3)},\dots,C_3^{(3)}$ as well. Let e be defined by $t_0,\dots,t_3 \in \mathbb{R}$ via the equation

$$(t_0,t_1,t_2)\cdot(x,y,z)+t_3 = 0.$$

We first claim that $t_0 \neq 0$. Indeed, if $t_0 = 0$, the projection of e into the yz-plane is a line. This line has to intersect the projections of the $C_i^{(3)}$. However, these projections give the subsets of the rectangles defined in Figure 7, and it can be seen with similar methods as in Theorem 6 that no such line exists for $\beta < 3$. Moreover, it can be easily verified that e is not parallel to any coordinate plane.

Since $t_0 \neq 0$, we can normalize such that $t_0 = 1$. We define the induced mapping

$$E : \mathbb{R}^3 \to \mathbb{R}, (x,y,z) \mapsto (1,t_1,t_2)\cdot(x,y,z)+t_3.$$

Clearly, $e = E^{-1}(0)$, and if for two points $p_1,p_2 \in \mathbb{R}^3$, $E(p_1)\cdot E(p_2) < 0$, then p_1 and p_2 are separated in different halfspaces by e. From this property, we can immediately deduce

LEMMA. e intersects each $C_i^{(3)}$, $1 \leq i \leq 4$ if and only if E satisfies the following four formulas

$$E((-1,-1,-1)) \cdot E((-3,-3,-3)) \leq 0$$

$$E((1,-1,-1)) \cdot E((3,-3,-3)) \leq 0$$

$$\begin{aligned} E((-1,1,-1)) \cdot E((3,3,-3)) < 0 \quad \vee \\ E((1,1,-1)) \cdot E((-3,3,-3)) < 0) \end{aligned}$$

$$\begin{aligned} E((-1,-1,1)) \cdot E((3,3,3)) < 0 \quad \vee \\ E((1,-1,1)) \cdot E((-3,3,3)) < 0 \quad \vee \\ E((-1,1,1)) \cdot E((3,-3,3)) < 0 \quad \vee \\ E((1,1,1)) \cdot E((-3,-3,3)) < 0 \end{aligned}$$

PROOF. The $\Leftarrow$-direction is simple to proof. For the $\Rightarrow$-direction, note that the first two statements are trivially satisfied. For the third statement, note first that e cannot contain the whole line segment from $(-1,1,-1)$ to $(1,1,-1)$, neither the line segment from $(-3,3,-3)$ to $(3,3,-3)$ (because this would contradict $t_0 \neq 0$). Let f be the plane that contains $C_2^{(3)}$. The intersection of e and f is a line on f that intersects the boundary of $C_2^{(3)}$ in exactly two points. It is therefore clear that this line separates two opposite vertices in two halfplanes. Thus, this pair is separated into two halfspaces by e and this is precisely what is checked by the third statement.

For the last statement, note that e cannot completely contain any boundary face of $C_3^{(3)}$, because this would mean that e is parallel to a coordinate plane which is easily checked to be impossible. Therefore, e intersects through the interior of $C_3^{(3)}$. Therefore e separates at least one of the four pairs of opposite vertices of the polytope, and this is checked in the forth statement. □

With Lemma 15, we can formulate a quantified formula in the real variables t_1, t_2, t_3 which is true if and only if a plane e as required exists. However, using a quantifier elimination program, we can easily compute that the formula is in fact false.

Singular case This case is completely analogous. For completeness, we write down the sets:

$$\begin{aligned} C_0^{(b)} &:= \mathrm{hull}\{(-1,-1,-1),(-b,-b,-b),(1,-1,-1),(b,-b,-b)\} \\ &\quad -\mathrm{hull}\{(-1,-1,-1),(-b,-b,-b)\} \\ &\quad -\mathrm{hull}\{(1,-1,-1),(b,-b,-b)\} \\ C_1^{(b)} &:= \mathrm{hull}\{(-1,1,-1),(-b,b,-b),(1,1,-1),(b,b,-b)\} \\ &\quad -\mathrm{hull}\{(-1,1,-1),(-b,b,-b)\} \\ &\quad -\mathrm{hull}\{(1,1,-1),(b,b,-b)\} \\ C_2^{(b)} &:= \mathrm{hull}\{(-1,-1,1),(-b,-b,b),(-1,1,1),(-b,b,b)\} \\ &\quad -\mathrm{hull}\{(-1,-1,1),(-b,-b,b)\} \\ &\quad -\mathrm{hull}\{(-1,1,1),(-b,b,b)\} \\ C_3^{(b)} &:= \mathrm{hull}\{(1,-1,1),(b,-b,b),(1,1,1),(b,b,b)\} \\ &\quad -\mathrm{hull}\{(1,-1,1),(b,-b,b)\} \\ &\quad -\mathrm{hull}\{(1,1,1),(b,b,b)\}. \end{aligned}$$

Defining e and E as above, we obtain:

LEMMA. e intersects each $C_i^{(3)}$, $1 \leq i \leq 4$ if and only if E satisfies the following four formulas

$$\begin{aligned} E((-1,-1,-1)) \cdot E((3,-3,-3)) < 0 \quad \vee \\ E((1,-1,-1)) \cdot E((-3,-3,-3)) < 0 \end{aligned}$$

$$\begin{aligned} E((-1,1,-1)) \cdot E((3,3,-3)) < 0 \quad \vee \\ E((1,1,-1)) \cdot E((-3,3,-3)) < 0 \end{aligned}$$

$$\begin{aligned} E((-1,-1,1)) \cdot E((-3,3,3)) < 0 \quad \vee \\ E((-1,1,1)) \cdot E((-3,-3,3)) < 0 \end{aligned}$$

$$\begin{aligned} E((1,-1,1)) \cdot E((3,3,3)) < 0 \quad \vee \\ E((1,1,1)) \cdot E((3,-3,3)) < 0. \end{aligned}$$

Again, it can be proven with a quantifier elimination program that the induced quantified formula is false.